Dr. Doug Cochran
Defense Sciences Office (DSO)
Reducing Complexity for Defense

Good Morning.

As Michael Goldblatt pointed out in his introductory remarks, a vigorous mathematics program is a key element in DSO. Michael, Steve Wax, Joe Bielietzki, and Valerie Browning have already mentioned a number of DSO activities in both biology and materials science where mathematics is playing an important enabling role. I want to tell you about some of our programs and initiatives in which mathematics is transforming science fiction into real defense technologies.

Let’s begin in the area of sensing. Without question, the performance of our sensors has improved dramatically in recent years. For example, the simple lead/tin-telluride hot spot photodetectors that were used in the first generation of Sidewinder seem downright primitive next to the large focal plane arrays in the latest generation of Sidewinder. We've gone from a single-pixel, mechanically scanned pattern to a high-resolution, rapidly refreshed image that the developers of the original Sidewinder could only dream about. From simple hard-wired radars that gave only range and bearing information, we have progressed to highly agile radar systems that provide high resolution images, detect moving targets, support multiple beams with adaptive beam patterns, and can change waveforms from pulse to pulse. It is not hard to imagine that the pioneers of radar at the MIT Rad Lab would think our modern systems to be pure science fiction.

But our vastly improved ability to gather data is not without its dark side: the curse of dimensionality. Even with ongoing improvements in computing, the availability of more and more data being produced by an ever- increasing number of sources can be overwhelming. The amount of data we produce has already swamped our image analysts. It is estimated that most of the data our sensors produce "hits the floor." It is never looked at or is only quickly scanned. In this respect, we are not doing very well against the dimensionality curse. And it's getting worse. Technology advances are leading to new sensors, based on both traditional and emerging modalities, that provide mountains of high-fidelity measurements of disparate types. For example, the emergence of hyperspectral sensors that monitor simultaneously 100 or more bands may raise computational burdens by two or more orders of magnitude. To make effective use of rapidly increasing amounts of available sensor data, we face the critical mathematical challenge of developing efficient ways to identify and extract what is truly important.

As we face this DARPA-hard challenge, there are two key weapons in our arsenal. The first is one we all know: the escalating power of computing hardware according to Moore's Law. But there are many important application areas in which Moore's law, even if indefinitely sustainable, may not be enough. Fortunately, we have a second tool: algorithms. Suppose our data rate increases by a factor of 100; can we really afford to wait 7 generations (about 10 years) before we can process the data in real time?

It's rather surprising to many people, but in most areas of science and engineering, with improvements in our ability to compute efficiently comes improved algorithms as well as hardware. In fact, I have found that even my mathematical colleagues are often unaware that algorithmic advances have accounted for about the same portion or more of our increase in data exploitation ability over the past 2 decades as have processor speed increases. Pushing forward advances in algorithm performance is a tenet of the DARPA Applied and Computational Mathematics Program. Central to this theme is incorporating as much of the physics as possible into our mathematical models so that crucial low-dimensional structure in the problem can be identified and exploited by the resulting algorithms.

In this spirit, the Integrated Sensing and Processing Program is creating methodologies and algorithms to intelligently select the data sensor systems actually collect. I will take a few moments to elaborate on this sensing program, which is moving us from a back-end-to-front-end feedback perspective toward a more holistic view of both sensor system operation and design. Current design methodologies for sensor systems generally create a feed-forward chain of subsystems that are functionally compatible, but are themselves independently designed and optimized—typically by separate communities of experts that don't talk much to each other. The resulting architectures pass information from one subsystem to the next, beginning with measurement and sampling processes on analog physical phenomenologies, through intermediate digital representations and processing, and to higher-level processing and exploitable symbolic information. Optimization of the individual subsystems in such a chain produces new capabilities in sensor system components, but is unlikely to produce the best possible global result. Moreover, it perpetuates existing bottlenecks of data flow and processing.

At the front end of the chain, improvements in individual sensor capabilities are sought. To be truly useful, they need to be matched to the job. For example, increased spatial resolution is not much help if you want to know if an object is purple. In fact, increasing the data rate without increasing its relevant information content unduly burdens the next link in the chain: the DSP algorithms that must refine, enhance, and integrate this raw digital data to facilitate discovery of structure hidden in raw sensor output and to facilitate its ultimate exploitation downstream. The bottleneck at this point presents a grave situation with regard to obtaining sufficient computational throughput for timely performance—even with continuing advances in hardware.

One way to deal effectively with the flood of data available from today's and tomorrow's sensor systems is to intelligently select the data one actually collects. Sensors used in reconnaissance, surveillance, and weapon systems are increasingly mode-switchable, tunable, and otherwise configurable. Real-time access to the operating parameters of sensors allows back-end exploitation algorithms, such as target classifiers and trackers, to control the front-end sensors. Imagine what would happen if our sensor systems worked this way. Only the data most needed would be collected. While simple in concept, this approach represents a significant departure from classical perspectives on signal detection, target classification, and tracking. Where algorithms for these functions have been designed to do the best they can with the data they are given.

The Integrated Sensing and Processing Program visualizes a future in which our sensors, in real time, determine what data they need to best perform their jobs. Imagine, for example, a waveform-agile radar whose software front end allows the next transmitted pulse to be designed "on the fly." Such a system could play a mathematically optimal game of "20 questions" with the scenario environment in which the selected waveforms play the role of the questions. But knowing which question will give the most useful information in view of what is already known, both from previous queries and ancillary information, presents significant new mathematical challenges that the Integrated Sensing and Processing Program is addressing. It is establishing a genuine scientific foundation for managing the transformation and flow of information among the functional elements of a sensor system. Innovative and effective mathematical optimization and control strategies form the cornerstone for intelligent integration of currently disjoint tasks of sensing and computation. This program will not only affect radar, but also sonar, ladar, and IR systems, as well as scheduling our sensors to work together cooperatively to finish the 20 questions game with only 5 questions.

The Integrated Sensing and Processing Program arises from a vision of how to compute well if the only tool available is a digital processor. Imagine new ways of computing based upon completely new foundations such as biocomputing and quantum computing that are inherently parallel, fast, or both. Stu Wolf, Mike Foster, and I are managing the Quantum Information Science and Technology Program, known as QuIST. Achieving success in quantum computing will require a marriage of advances in physics, computational science, and control mathematics. Exploiting quantum computers will require new and strange ways of looking at algorithms. One key mathematical challenge is in control of quantum systems. It is a guiding principle of quantum mechanics that in order to measure a state you must destroy it. But, to exploit a quantum state in computing, you must control it. Within the QuIST Program, control mathematicians are exploring the delicate problem of controlling quantum states without destroying them. In the vein of quantum algorithms, researchers have recently factored the number 15. While this may not sound very hard, it is, in fact, a breakthrough that signals the beginning of a new era in which the ability to factor large numbers will change the face of cryptography. We anticipate being able to perform calculations that are beyond our abilities today. And as with conventional computers, we will put quantum computers to use in ways we have not yet begun to imagine. Fifteen today; 100-digit numbers in a few years?

Let me return to the theme of how to compute well. In recent years, advances in computation have been coupled with increasingly sophisticated techniques for modeling and simulation to develop virtual design capabilities. The primary objective of virtual design is to reduce design cycle time and cost by carrying forward as much of the process as possible within a computer-based modeling and simulation environment. If we have the ability to rapidly compute models that agree with the experiment, we can replace costly build, test, and evaluate cycles, with relatively cheap simulations performed on computers. Virtual design is already standard practice in several engineering disciplines. Perhaps the most successful example is SPICE, a widely-used software package that allows electronic components to be designed and tested entirely within a computer environment before undertaking the costly and time-consuming fabrication process.

Development of such an environment requires that the physics underlying design criteria be captured in a mathematical model that provides both high fidelity and computational tractability. SPICE got its start with physical models accurately described by relatively simple and well-understood mathematical machinery, such as analysis of simple differential equations. Continuing developments in virtual design of electronic circuits and other products has led to incorporation of increasingly sophisticated mathematical representations of the products and substantially more complicated mathematical models to support simulation of performance. In many application arenas, computational limitations are the main factor restricting the scope of virtual design. As a consequence, reducing the complexity of models and algorithms has become a vital concern.

Let's ask a DARPA-hard question: What would it take to do a virtual design of a full-size aircraft? Today, the computational fluid dynamics required for such a design is almost routine, except perhaps for special cases such as, say, morphing materials. The same cannot be said of our ability to predict the radar cross-section associated with the design. Today, we model small sections of an aircraft independently. We build subscale and full-scale physical models and perform expensive tests on ranges. Our systems are costly and take a long time to design. And, worse, we don't even know if have an optimal design when we finish.

Now imagine a world in which an engineer sits down at a workstation and from a single CAD file is able to simulate not only the aerodynamic properties of an aircraft, but its electromagnetic scattering properties as well. That is precisely the vision of our Virtual Electromagnetic Testrange Program. This program is targeting dramatic reductions in the cost and time associated with accurate predictions of radar cross-section for full-size air vehicles with realistic material treatments and component details including cavities, thin edges, and embedded antennas. This program has as its foundation new classes of numerical analytic methods, developed under DARPA sponsorship and that of other DoD sponsors. A key breakthrough that we are exploiting in this program is a new algorithm, called the Fast Multipole Method, for the solution of integral equations that arise in scattering theory. In the very recent past, we were able to calculate the radar cross-section of the aircraft in isolation at moderate frequencies. Now we are able to calculate the complete aircraft radar cross-section, primarily as a result of an improved algorithm. As computers advance, we will be able to calculate cross-sections at higher frequencies for free.

The Fast Multipole Method exemplifies the kinds of complexity reduction attainable from algorithm enhancement; it decreases the computational cost for a broad class of scattering calculations from order n2 to order n log(n), where *n* represents the number of nodes in the discretization of the domain of the integral operator. To put this in perspective, the impact of the new algorithm is the equivalent of skipping on the order of 15 generations of Moore's Law!

In our future world, we've moved the computational time for the radar cross-section of a full aircraft to a few hours. Now imagine that we've applied these capabilities to calculate the cross-section of a target in ground clutter or a naval vessel on the open sea. Such abilities will help us identify features that lead to enhanced performance of our ATR systems. This capability will certainly change the way we design our future systems and aid in the discovery of new geometries and material treatments that are even stealthier than today's.

What comes to mind when I ask you to think about a map? Some of you may think of roadmaps, some might think of lines of isoaltitude printed on a paper, some might think of a matrix whose entries are altitude values as in a digital elevation model. All these are representations of data. How we represent data affects the way in which we use data and the efficiency with which we do so. The Computational and Algorithmic Representation of Geometric Objects Program, conducted jointly with the National Science Foundation, imagines a future in which we are able to compactly and efficiently organize geospatial data. From point cloud data, we are able to obtain objects. From isolated points, such as in a DEM, we rapidly extract line-of-sight information while having efficiently compressed the underlying data by several orders of magnitude. Representation is not only important for geospatial data, but also in biology where an efficient representations can help us understand protein folding and make progress toward rapid synthesis of agents to counteract chemical and biological threats. Imagine the impact of being able to host on your PDA a highly detailed map that not only lets you know where you are, but which stores are nearby, and their entire current inventories and price lists. This is just another example of turning science fiction into science fact through the exercise of imagination.

My focus today has been on sharing with you the increasingly important role of complexity reduction in representations and algorithms for defense applications. We are exploring new perspectives on sensor system design and operation, where a marriage of technological and mathematical advances holds great promise for development of integrated sensing and processing architectures and operational concepts. Important challenges for complexity reduction also remain in a spectrum of virtual design applications, including air and sea platforms, electronic and structural materials, and antennas.

New opportunities and challenges continue to emerge, especially in connection with the recent explosion of quantitative perspectives across many aspects of biology. We are beginning to imagine a world in which we will be able to perform virtual design for pharmaceuticals and design materials that do not occur in nature for a wide variety of applications. We can even conceive that a brain-machine interface based on integrated sensing and processing principles will come about in the next decade. Our legacy will entail not only the impact of our imagination on a number of fields, but a set of tools that can interact with the imagination to change the world in ways that we have yet to conceive.

I thank you for your attention and the opportunity to share with you a vision of how mathematics will shape our future.